

Stat 134 Fall 2011: Gambler's ruin

Michael Lugo

September 12, 2011

In class today I talked about the problem of “gambler’s ruin” but there wasn’t enough time to do it properly. I fear I may have confused some people by not stating some things as clearly as I could have, in the interest of saving time. Here’s what I meant to say.

Alice and Bob are betting on the outcomes of the flip of a coin, which comes up heads with probability p , and tails with probability $q = 1 - p$. Alice starts out with a units, and Bob starts out with b units. (In class I used r and s but I want those letters later.) We have $n = a + b$. They flip the coin repeatedly; each time it comes up heads, Bob pays Alice 1 unit, and each time it comes up tails, Alice pays Bob 1 unit. The game ends when one player has no money; the other player has n units at that time, and we say they win.

There are two natural questions to ask:

- What’s the probability that Alice wins? (Alternatively, what’s the probability that Bob wins? It turns out that this game will, with probability 1, eventually terminate with one player winning or the other, even though you could imagine it going on forever.)
- How long do we expect this game to last?

The first question is something we can answer now; the second question we can’t, mostly because we don’t know what “expect” means. Maybe we’ll revisit this later in the semester once we’ve talked about expected value.

Some special cases

The probability of Alice winning is of course a function of a, n and p . (We could say it’s a function of a, b, p or of b, n, p , but we won’t.) We’ll denote this function by $f(a, n, p)$.

For some special cases we can find the probabilities by path counting.

a = 1, n = 2. Alice has one unit, and Bob has one unit. So Alice wins if the first toss comes up heads, and loses if it comes up tails. Therefore $f(1, 2, p) = p$.

a = 1, n = 3. Alice has one unit, and Bob has two units. Therefore:

- if the first toss is a tail, Alice loses. This happens with probability q .
- if the first two tosses are both heads, Alice wins. This happens with probability p^2 .

- if the first three tosses are H, T, T , Alice loses. This happens with probability pqq .
- if the first four tosses are H, T, H, H , Alice wins. This happens with probability pqp^2 .

The pattern continues, with Alice winning if the first $2k$ tosses, for some k , are $k - 1$ repetitions of H, T , followed by H, H . This happens with probability $(pq)^{k-1}p^2$. Therefore Alice wins with probability

$$p^2 + pqp^2 + (pq)^2p^2 + \dots = p^2 (1 + pq + (pq)^2 + \dots) = \frac{p^2}{1 - pq}$$

and so $f(1, 3, p) = p^2/(1 - pq)$.

a = 2, n = 3. In the previous case, interchange the pairs “Alice” and “Bob”; p and q ; “heads” and “tails”.

When Alice starts with one unit, and Bob starts with two units, and heads have probability p , Alice has probability $p^2/(1 - pq)$ of winning.

So when Bob starts with one unit, and Alice starts with two units, and tails have probability q , Bob has probability $q^2/(1 - qp)$ of winning.

In this case Alice has probability $1 - q^2/(1 - qp)$ of winning. Simplifying, this gives $f(2, 3, p) = p/(1 - pq)$.

a = 2, n = 4. Alice has two units, and Bob has two units. Therefore:

- if the first two tosses are heads, Alice wins. This happens with probability p^2 .
- if the first two tosses are tails, Alice loses. This happens with probability q^2 .
- if the first two tosses are a head and a tail, in either order, then we’re back where we started. This happens with probability $2pq$.

So the probability that the game will return to the starting position k times before Alice finally wins is $(2pq)^k p^2$. The probability that Alice wins is therefore

$$\sum_{k=0}^{\infty} (2pq)^k p^2 = \frac{p^2}{1 - 2pq} = \frac{p^2}{p^2 + q^2}.$$

and this is $f(2, 4, p)$.

a = 1, n = 4. Alice has one unit, Bob has three. If the first toss is a head then we’re in the situation of the previous case. If it’s a tail Alice loses right away. So

$$f(1, 4, p) = pf(2, 4, p) = \frac{p^3}{p^2 + q^2}.$$

a = 3, n = 4. Alice has three units, Bob has one. If the first toss is a head Alice wins. If it’s a tail then we’re in the case $a = 2, n = 4$. So

$$f(3, 4, p) = p + (1 - p)f(2, 4, p) = \frac{p(p^2 + q)}{p^2 + q^2}.$$

n = 2a, p = 1/2. Alice and Bob start out with the same amount of money. The coin is fair. It seems that the game is fair, and we guess $f(a, 2a, 1/2) = 1/2$. This isn’t quite a proof, though.

The general case

Consider $n = 5$; this is when it gets quite a bit harder to come up with arguments like the ones in the previous section that are based on counting the ways the game could evolve. So let's take a more probabilistic approach. Let r_k be the probability that Alice wins when she starts with k , and Bob starts with $n - k$, and the probability of heads is p . So $r_k = f(k, n, p)$, but it'll be easier to write r_k .

If $k = 0$ or $k = n$ the game is over before it starts, with Alice losing and winning, respectively; so $r_0 = 0, r_n = 1$.

Otherwise we can break down what happens into two cases. Either the first toss is a head – and then Alice has $k + 1$ – or the first toss is a tail – and then Alice has $k - 1$. So we have the recurrence

$$r_k = pr_{k+1} + qr_{k-1}.$$

You can check that the formulas we already derived satisfy this. But now we need some algebraic trickery to solve this recurrence in general. Write $r_k = pr_k + qr_k$ to get

$$pr_k + qr_k = pr_{k+1} + qr_{k-1}.$$

Subtract from both sides to get

$$pr_{k+1} - pr_k = qr_k + qr_{k-1}.$$

We now let $s_k = r_k - r_{k-1}$ be the difference between two consecutive r_k . Then we have

$$ps_{k+1} = qs_k$$

and solving for s_{k+1} gives $s_{k+1} = (q/p)s_k$. Therefore we have

$$s_2 = \frac{q}{p}s_1, s_3 = \frac{q}{p}s_2 = \left(\frac{q}{p}\right)^2 s_1, \dots, s_k = \left(\frac{q}{p}\right)^{k-1} s_1.$$

Now notice $r_1 = s_1, r_2 = s_1 + s_2$, and so on. We get the formula

$$r_k = s_1 + s_2 + \dots + s_k = s_1 \left(1 + \frac{q}{p} + \dots + \left(\frac{q}{p}\right)^{k-1} \right) = s_1 \frac{1 - (q/p)^k}{1 - (q/p)}. \quad (1)$$

But we know that

$$r_n = s_1 + s_2 + \dots + s_n = (r_1 - r_0) + (r_2 - r_1) + \dots + (r_n - r_{n-1}) = r_n - r_0 = 1 - 0 = 1.$$

In particular we can use equation (1) with $k = n$ to get

$$r_n = s_1 \frac{1 - (q/p)^n}{1 - (q/p)}$$

and we know $r_n = 1$, so we have

$$s_1 = \frac{1 - q/p}{1 - (q/p)^n}.$$

Finally, then we get

$$r_k = s_1 \frac{1 - (q/p)^k}{1 - (q/p)^n} = \frac{1 - (q/p)^k}{1 - (q/p)^n}$$

and restoring the original notation gives the formula

$$f(k, n, p) = \frac{1 - (q/p)^k}{1 - (q/p)^n}$$

for the probability that Alice wins starting with k units, when Bob starts with $n - k$ units, and the probability of heads is p .

If $p = 1/2$, then $q/p = 1$ and we can't use this formula. But in that case you get $s_{k+1} = s_k$ – that is, the differences between consecutive r_k don't depend on k – and so $f(k, n, 1/2) = k/n$.

Who cares?

You are in Las Vegas, and it will cost you 200 dollars to get back to Berkeley. You have 100 dollars. You are going to try to play roulette, betting on red at each spin of the wheel, in order to make enough money to get back. This is like flipping coins, except $p = 18/38$. How should you bet?

Let's say you're going to bet $100/x$ dollars at each spin of the wheel, where x is an integer. Then we'll let the "unit" of money be $100/x$ dollars. You have x units and you want $2x$ units; the probability that you'll get your 200 dollars back before you run out of money is $f(x, 2x, p)$. Using the formula we already derived,

$$f(x, 2x, p) = \frac{1 - (10/9)^x}{1 - (10/9)^{2x}}$$

since $q/p = (1 - p)/p = (20/38)/(18/38) = 10/9$. We can rewrite this as

$$f(x, 2x, p) = \frac{(10/9)^x - 1}{(10/9)^{2x} - 1} = \frac{1 - (9/10)^x}{(10/9)^x - (9/10)^x}.$$

This seems like a silly way to rewrite things, but actually it's not. If x is large, we can ignore the terms $(9/10)^x$ and so we get $f(x, 2x, p) \approx 1/(10/9)^x = (9/10)^x$. So if your bets are small you have a very small probability of winning. In terms of limits, $\lim_{x \rightarrow \infty} f(x, 2x, p) = 0$ when $p < 1/2$ – as the size of the bet gets small relative to the bankroll, the probability of winning approaches zero.

Furthermore, it's easy to see (by graphing) that $f(x, 2x, p)$ is a decreasing function of x . As x gets bigger – that is, as your bets get smaller, the probability of winning gets smaller.

The optimal strategy is to bet everything on a single spin – that is, to let $x = 1$ – in which case your probability of winning is $f(1, 2, 18/38) = 18/38$.

But now look at this from the point of view of the casino. From the casino’s point of view $p = 20/38$ – they win exactly when you lose. So the probability that the casino wins becomes *larger* as the ratio (total amount of money)/(single bet) – the number of “units” of money at risk – gets larger. You can show that $\lim_{x \rightarrow \infty} f(x, 2x, p) = 1$ if $p > 1/2$. The house always wins.

(Even if you had $p = 1/2$, in reality the casino has much more money than you. We have, for example, $f(100, 10^6, 1/2) = 100/10^6 = 10^{-4}$ – so if you show up in Vegas with one hundred dollars to gamble, and the casino has a million dollars, and you are offered fair games, you have one chance in ten thousand of bankrupting the casino. Of course the actual chance that a single gambler will bankrupt the casino is much less than this; it must be since more than ten thousand people gamble in Vegas every day. I just made this fact up, but it seems right.)

The moral of the story is as follows. If the odds are against you, and you need to get to a certain amount, bet big – this means that you have to make *less* bets, and the law of large numbers (which we’ll study more later) doesn’t have time to hurt you. If the odds are in your favor, bet small – then you make more bets, and the law of large numbers is likely to work in your favor.